What is claimed is:

1. A method of determining a placement of services of a distributed application onto nodes of a distributed resource infrastructure comprising the steps of:
   a. forming communication constraints between node pairs which ensure that a sum of transport demands between a particular node pair does not exceed a transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables;
   b. forming an objective; and
   c. employing a local search solution to solve an integer program comprising the communication constraints and the objective, which determines the placement of the services onto the nodes.

2. A method of determining a placement of services of a distributed application onto nodes of a distributed resource infrastructure comprising the steps of:
   a. establishing an application model of the services comprising transport demands between the services;
   b. establishing an infrastructure model of the nodes comprising transport capacities between the nodes;
   c. forming an integer program that comprises:
      i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;
      ii. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables; and
      iii. an objective; and
   d. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.

3. The method of claim 2 wherein the step of solving the integer program employs a local search solution.

4. The method of claim 2 wherein the objective comprises minimizing communication traffic between the nodes.

5. The method of claim 2 wherein the application model further comprises processing demands for the services.

6. The method of claim 5 wherein the infrastructure model further comprises processing capacities for the nodes.

7. The method of claim 6 wherein the integer program further comprises processing constraints which ensure that a sum of the processing demands for each of the nodes does not exceed the processing capacity for the node.

8. The method of claim 7 wherein the objective comprises minimizing communication traffic between the nodes and balancing the processing demands on the nodes.

9. The method of claim 6 wherein the processing demands and the processing capacities are normalized according to a processing criterion.

10. The method of claim 9 wherein the processing criterion comprises an algorithm speed.

11. The method of claim 9 wherein the processing criterion comprises a transaction speed.

12. The method of claim 9 wherein the processing capacities of the nodes are found according to a look-up table in which different types of nodes have been normalized according to the processing criterion.

13. The method of claim 2 wherein the application model further comprises storage demands for the services.

14. The method of claim 13 wherein the infrastructure model further comprises storage capacities for the nodes.

15. The method of claim 14 wherein the integer program further comprises storage constraints which ensure that a sum of the storage demands for each of the nodes does not exceed the storage capacity for the node.

16. The method of claim 2 wherein the integer program further comprises placement constraints which ensure that each of the services is placed on one and only one of the nodes.

17. The method of claim 2 wherein the services reside on the nodes according to a previous assignment.

18. The method of claim 17 further comprising the step of assessing reassignment penalties for service placements that differs from the previous assignment.

19. The method of claim 18 wherein the integer program further comprises a second objective that seeks to minimize the reassignment penalties.

20. A method of determining a placement of services of a distributed application onto nodes of a distributed resource infrastructure comprising the steps of:

a. establishing an application model of the services that comprises processing demands for the services, storage demands for the services, and transport demands between the services;

b. establishing an infrastructure model of the nodes that comprises processing capacities for the nodes, storage capacities for the nodes, and transport capacities between the nodes;

c. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular

service is located on a particular node;

ii. processing constraints which ensure that a sum of the processing demands for each of the nodes does not exceed the processing capacity for the node;

iii. storage constraints which ensure that a sum of the storage demands for each of the nodes does not exceed the storage capacity for the node;

iv. placement constraints which ensure that each of the services is placed on one and only one node;

v. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables; and

vi. an objective of minimizing communication traffic between the nodes and balancing processing loads on the nodes; and

d. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.

21. A computer readable memory comprising computer code for directing a computer to make a determination of a placement of services of a distributed application onto nodes of a distributed resource infrastructure, the determination of the placement of the services onto the nodes comprising the steps of:

a. forming communication constraints between node pairs which ensure that a sum of transport demands between a particular node pair does not exceed a transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables;

b. forming an objective; and

c. employing a local search solution to solve an integer program comprising the communication constraints and the objective, which determines the placement of the services onto the nodes.

22. A computer readable memory comprising computer code for directing a computer to make a determination of a placement of services of a distributed application onto nodes of a distributed resource infrastructure, the determination of the placement of the services onto the nodes comprising the steps of:

a. establishing an application model of the services comprising transport demands between the services;

b. establishing an infrastructure model of the nodes comprising transport capacities between the nodes;

c. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;

ii. communication constraints between node pairs which ensure that a sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables; and

iii. an objective; and

d. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.

23. The computer readable memory of claim 22 wherein the step of solving the integer program employs a local search solution.

24. The computer readable memory of claim 22 wherein the objective comprises minimizing communication traffic between the nodes.

25. The computer readable memory of claim 22 wherein the application model further comprises processing demands for the services.

26. The computer readable memory of claim 25 wherein the infrastructure model further comprises processing capacities for the nodes.

27. The computer readable memory of claim 26 wherein the integer program further comprises processing constraints ensure that a sum of the processing demands for each of the nodes does not exceed the processing capacity for the node.

28. The computer readable memory of claim 27 wherein the objective comprises balancing the processing demands on the nodes.

29. The computer readable memory of claim 26 wherein the processing demands and the processing capacities are normalized according to a processing criterion.

30. The computer readable memory of claim 29 wherein the processing criterion comprises an algorithm speed.

31. The computer readable memory of claim 9 wherein the processing criterion comprises a transaction speed.

32. The computer readable memory of claim 9 wherein the processing capacities of the nodes are found according to a look-up table in which different types of nodes have been normalized according to the processing criterion.

33. The computer readable memory of claim 22 wherein the application model further comprises storage demands for the services.

34. The computer readable memory of claim 33 wherein the infrastructure model further comprises storage capacities for the nodes.

35. The computer readable memory of claim 34 wherein the integer program further comprises storage constraints which ensure that a sum of the storage demands for each of the nodes does not exceed the storage capacity for the node.

36. The computer readable memory of claim 22 wherein the integer program further comprises placement constraints which ensure that each of the services is placed on one and only one of the nodes.

37. The computer readable memory of claim 22 wherein the services reside on the nodes according to a previous assignment.

38. The computer readable memory of claim 37 further comprising the step of assessing reassignment penalties for service placements that differs from the previous assignment.

39. The computer readable memory of claim 38 wherein the integer program further comprises a second objective that seeks to minimize the reassignment penalties.

40. A computer readable memory comprising computer code for directing a computer to make a determination of a placement of services of a distributed application onto nodes of a distributed resource infrastructure, the determination of the placement of the services onto the nodes comprising the steps of:

a. establishing an application model of the services that comprises processing demands for the services, storage demands for the services, and transport demands between the services;

b. establishing an infrastructure model of the nodes that comprises processing capacities for the nodes, storage capacities for the nodes, and transport capacities between the nodes;

c. forming an integer program that comprises:

i. a set of placement variables for a combination of the services and the nodes, each of the placement variables indicating whether a particular service is located on a particular node;

ii. processing constraints which ensure that a sum of the processing demands for each of the nodes does not exceed the processing capacity for the node;

iii. storage constraints which ensure that a sum of the storage demands for each of the nodes does not exceed the storage capacity for the node;

iv. placement constraints which ensure that each of the services is placed on one and only one node;

v. communication constraints between node pairs which ensure that a

sum of the transport demands between a particular node pair does not exceed the transport capacity between the particular node pair, each term of the sum comprising a product of a first placement variable, a second placement variable, and the transport demand between the services associated with the first and second placement variables; and

vi. an objective of minimizing communication traffic between the nodes and balancing processing loads on the nodes; and

d. employing a local search solution to solve the integer program which determines the placement of the services onto the nodes.